# EAST Search History

| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|---|---|---|---|---|---|---|
| L1 | 1 | rank$4 same module$1 same quer$4 same (database$1 (data adj base$1) db$1) same rang$4 same deviat$4 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/09/12 11:41 |
| L2 | 1 | sort$4 same modul$4 same rank$4 same (database$1 db$1 (data adj base$1)) same rang$4 same deviat$4 same order$4 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/09/12 11:42 |

Terms used
**rank$4 same module$1 same quer$4 same database$1 data adj base$1 db$1 same rang$4 same deviat$4**

Found **107,772** of **185,030**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

**1** Shape-based retrieval and analysis of 3D models

Thomas Funkhouser, Michael Kazhdan

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes SIGGRAPH '04**

**Publisher:** ACM Press

Full text available: pdf(12.56 MB) Additional Information: full citation, abstract

Large repositories of 3D data are rapidly becoming available in several fields, including mechanical CAD, molecular biology, and computer graphics. As the number of 3D models grows, there is an increasing need for computer algorithms to help people find the interesting ones and discover relationships between them. Unfortunately, traditional text-based search techniques are not always effective for 3D models, especially when queries are geometric in nature (e.g., find me objects that fit into thi ...

**2** Research papers: data cleaning and mapping: A cost-based model and effective heuristic for repairing constraints by value modification

Philip Bohannon, Wenfei Fan, Michael Flaster, Rajeev Rastogi

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

**Publisher:** ACM Press

Full text available: pdf(565.83 KB) Additional Information: full citation, abstract, references

Data integrated from multiple sources may contain inconsistencies that violate integrity constraints. The *constraint repair problem* attempts to find "low cost" changes that, when applied, will cause the constraints to be satisfied. While in most previous work repair cost is stated in terms of tuple insertions and deletions, we follow recent work to define a database repair as a set of *value modifications.* In this context, we introduce a novel cost framework that allows for the appl ...

**3** Special issue: AI in engineering

D. Sriram, R. Joobbani

April 1985 **ACM SIGART Bulletin**, Issue 92

**Publisher:** ACM Press

Full text available: pdf(8.79 MB) Additional Information: full citation, abstract

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

4 Consistency and orderability: semantics-based correctness criteria for databases

Divyakant Agrawal, Amr El Abbadi, Ambuj K. Singh

September 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 3

**Publisher:** ACM Press

Full text available: pdf(1.92 MB) Additional Information: full citation, abstract, references, citings, index terms

The semantics of objects and transactions in database systems are investigated. User-defined predicates called consistency assertions are used to specify user programs. Three new correctness criteria are proposed. The first correctness criterion consistency is based solely on the users' specifications and admit nonserializable executions that are acceptable to the users. Integrity constraints of the database are maintained through consistency assertions. Th ...

**Keywords:** concurrency control, object-oriented databases, semantics, serializability theory

5 Concurrency control: methods, performance, and analysis

Alexander Thomasian

March 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 1

**Publisher:** ACM Press

Full text available: pdf(427.18 KB) Additional Information: full citation, references, citings, index terms

**Keywords:** Markov chains, adaptive methods, concurrency control, data contention, deadlocks, flow diagrams, load control, optimistic concurrency control, queueing network models, restart-oriented locking methods, serialiazability, thrashing, two-phase locking, two-phase processing, wait depth limited methods

6 A prototype implementation of the SQL Ada module extension (SAME) method

Allison LeClair, Susan Phillips

December 1990 **Proceedings of the conference on TRI-ADA '90**

**Publisher:** ACM Press

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings

As Ada becomes more widespread, the ability to access commercial database technologies through Ada systems becomes a significant issue. Researchers throughout our industry are investigating interface approaches between Ada and these technologies, including language bindings between Ada and SQL, a relational data base language. This paper presents a recent implementation of one such binding—the SQL Ada Module Extension (SAME) method.

7 A distributed object-oriented database system supporting shared and private databases

Won Kim, Nat Ballou, Jorge F. Garza, Darrell Woelk

January 1991 **ACM Transactions on Information Systems (TOIS)**, Volume 9 Issue 1

**Publisher:** ACM Press

Full text available: pdf(1.58 MB) Additional Information: full citation, abstract, references, citings, index terms, review

ORION-2 is a commercially available, federated, object-oriented database management system designed and implemented at MCC. One major architectural innovation in ORION-2 is the coexistence of a shared databese and a number of private databases. The shared database is accessible to all authorized users of the system, while each private database is accessible to only the user who owns it. A distributed database system with a shared database and private databases for individual users is a natu ...

**Keywords:** client-server architecture, federated databases, object-oriented databases

8 Pervasive Documentation Systems I: Concept and architecture of an pervasive document editing and managing system
Stefania Leone, Thomas B. Hodel, Harald Gall
September 2005 **Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information SIGDOC '05**
**Publisher:** ACM Press
Full text available: pdf(489.38 KB) Additional Information: full citation, abstract, references, index terms

Collaborative document processing has been addressed by many approaches so far, most of which focus on document versioning and collaborative editing. We address this issue from a different angle and describe the concept and architecture of a pervasive document editing and managing system. It exploits database techniques and real-time updating for sophisticated collaboration scenarios on multiple devices. Each user is always served with up-to-date documents and can organize his work based on docu ...

**Keywords:** collaborative document, computer supported collaborative work (CSCW), pervasive document editing and management system

9 Soviet Computer Technology—1959
March 1960 **Communications of the ACM**, Volume 3 Issue 3
**Publisher:** ACM Press
Full text available: pdf(8.23 MB) Additional Information: full citation

10 Scalable packet classification
Florin Baboescu, George Varghese
August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '01**, Volume 31 Issue 4
**Publisher:** ACM Press
Full text available: pdf(242.61 KB) Additional Information: full citation, abstract, references, citings, index terms

Packet classification is important for applications such as firewalls, intrusion detection, and differentiated services. Existing algorithms for packet classification reported in the literature scale poorly in either time or space as filter databases grow in size. Hardware solutions such as TCAMs do not scale to large classifiers. However, even for large classifiers (say 100,000 rules), any packet is likely to match a few (say 10) rules. Our paper seeks to exploit this observation to produce a s ...

11 Outerjoin optimization in multidatabase systems
Arbee L. P. Chen
July 1990 **Proceedings of the second international symposium on Databases in parallel and distributed systems**
**Publisher:** ACM Press
Full text available: pdf(802.60 KB) Additional Information: full citation, abstract, references, citings, index terms

Outerjoin is used in distributed relational multidatabase systems for integrating local schemas to a global schema. Queries against the global schema need to be modified, optimized, and decomposed into subqueries at local sites for processing. Since outerjoin combines local relations in different databases to form a global relation, it is expensive to process. In this paper, based on the structure of the query and the definition of the schemas, queries with outerjoin, join, select and proje ...

12 A Data-Model Interface for modular dynamic simulation
Friedhelm Drepper
January 1978

**Proceedings of the 10th conference on Winter simulation - Volume 1**

**Publisher:** IEEE Press

Full text available: pdf(497.17 KB) Additional Information: full citation, abstract, references, index terms

As an alternative to simulation languages like Dynamo a Fortran based simulation aid is presented which offers more flexibility at the cost of some conceptual guidance. This tool is essentially a Data-Model Interface (DMI). It supports a modular model structure and is especially suited for the simulation of large dynamic systems with a simple event structure. Its simulation concept allows unsorted statements as well as coupled equation systems extending over several modules.

**13** Designing interconnection networks for multi-level packaging

M. T. Raghunath, A. Ranade

December 1993 **Proceedings of the 1993 ACM/IEEE conference on Supercomputing**

**Publisher:** ACM Press

Full text available: pdf(1.12 MB) Additional Information: full citation, references, citings, index terms

**14** Federated databases and systems: part II --- a tutorial on their resource consolidation

David K. Hsiao

October 1992 **The VLDB Journal — The International Journal on Very Large Data Bases,** Volume 1 Issue 2

**Publisher:** Springer-Verlag New York, Inc.

Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references

The issues and solutions for the interoperability of a class of heterogeneous databases and their database systems are expounded in two parts. Part I presented the data-sharing issues in federated databases and systems (Hsiao, 1992). The present article explores resource-consolidation issues. *Interoperability* in this context refers to data sharing among heterogeneous databases, and to resource consolidation of computer hardware, system software, and support personnel. *Resource consolid ...*

***Keywords:*** *attribute-based, data-model-and-language-to-data-model-and-language mappings, database conversion, hierarchical, network, object-oriented, relational, schema transformation, transaction translation*

**15** Natural language information retrieval in digital libraries

Tomek Strzalkowski, Jose Perez-Carballo, Mihnea Marinescu

April 1996 **Proceedings of the first ACM international conference on Digital libraries**

**Publisher:** ACM Press

Full text available: pdf(1.03 MB) Additional Information: full citation, references, index terms

**16** An approach for designing decision support systems

Eric D. Carlson

December 1978 **ACM SIGMIS Database,** Volume 10 Issue 3

**Publisher:** ACM Press

Full text available: pdf(1.09 MB) Additional Information: full citation, abstract, references

Studies of specific decisions and general studies of decision making have indicated the potential benefits of computer support for decision making. These potential benefits can be divided into two categories: displaced cost and added value.

**17** Parallel evaluation of multi-join queries

Annita N. Wilschut, Jan Flokstra, Peter M. G. Apers

May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international**

**conference on Management of data SIGMOD '95**, Volume 24 Issue 2
**Publisher:** ACM Press
Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references, citings, index terms

A number of execution strategies for parallel evaluation of multi-join queries have been proposed in the literature; their performance was evaluated by simulation. In this paper we give a comparative performance evaluation of four execution strategies by implementing all of them on the same parallel database system, PRISMA/DB. Experiments have been done up to 80 processors. The basic strategy is to first determine an execution schedule with minimum total cost and then parallelize this schedule w ...

**18** Query Optimization: Predicting the cost-quality trade-off for information retrieval queries: facilitating database design and query optimization
Henk Ernst Blok, Djoerd Hiemstra, Sunil Choenni, Franciska de Jong, Henk M. Blanken, Peter M.G. Apers
October 2001 **Proceedings of the tenth international conference on Information and knowledge management**
**Publisher:** ACM Press
Full text available: pdf(1.42 MB) Additional Information: full citation, abstract, references, citings, index terms

Efficient, flexible, and scalable integration of full text information retrieval (IR) in a DBMS is not a trivial case. This holds in particular for query optimization in such a context. To facilitate the bulk-oriented behavior of database query processing, a priori knowledge of how to limit the data efficiently prior to query evaluation is very valuable at optimization time. The usually imprecise nature of IR querying provides an extra opportunity to limit the data by a trade-off with the qualit ...

**Keywords:** Zipf, databases, efficiency, fragmentation, information retrieval, quality, trade-off

**19** Constructive deterministic PRAM simulation on a mesh-connected computer
Andrea Pietracaprina, Geppino Pucci, Jop F. Sibeyn
August 1994 **Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures**
**Publisher:** ACM Press
Full text available: pdf(930.49 KB) Additional Information: full citation, abstract, references, index terms, review

We present a constructive deterministic simulation of a PRAM with n processors and m = n &agr; shared variables, 1 < &agr; ⩽ 2, on an n-node mesh-connected computer where each node hosts a processor and a memory module. At the core of the simulation is a Hierarchical Memory Organization Scheme (HMOS) that governs the distribution of the PRAM variables (each replicated into a number of copies) amon ...

**20** DB-1 (databases): data integration: Organizing structured web sources by query schemas: a clustering approach
Bin He, Tao Tao, Kevin Chen-Chuan Chang
November 2004 **Proceedings of the thirteenth ACM international conference on Information and knowledge management CIKM '04**
**Publisher:** ACM Press
Full text available: pdf(323.72 KB) Additional Information: full citation, abstract, references, index terms, review

In the recent years, the Web has been rapidly "deepened" with the prevalence of databases online. On this deep Web, many sources are structured by providing structured query interfaces and results. Organizing such structured sources into a domain hierarchy is one of the critical steps toward the integration of heterogeneous Web sources. We observe that, for structured Web sources, query schemas ie, attributes in query interfaces) are discriminative representative ...

**Keywords:** data integration, deep Web, hierarchical agglomerative clustering

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next